International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 04 | April - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Algorithmic Accountability in Sustainable Finance: A Systematic Review of AI, NLP, and Machine Learning Integration in Global ESG Risk Evaluation

Siddharth Shetty, Dhruvi Patel, Khushi Patil, Shivanand Pandey, Chaitali Mhatre

Department of Computer Engineering, Universal College of Engineering, Mumbai, India

siddharthshetty1982@gmail.com | dhruvipatel201004@gmail.com | khushipatil220104@gmail.com | pandeyshivanand753@gmail.com | patilchaitali1333@gmail.com

Abstract—Over the past decade, ESG criteria have moved from the periphery of investment analysis to its centre, yet the infrastructure supporting ESG evaluation has not kept pace. Rating agencies disagree with each other, companies report what suits them, and the smaller businesses that make up the vast majority of the global economy are effectively shut out of sustainable finance markets entirely. This paper examines how AI, Natural Language Processing (NLP), and machine learning are being deployed to address these gaps, drawing on a systematic review of eleven empirical studies. The findings are, in places, counterintuitive. Simple Linear Regression outperforms LSTM networks and BERT on corporate ESG text corpora — achieving 68.09% accuracy compared to 25.53% and 46.81% respectively — because corporate sustainability reports are so formulaic that deep learning models overfit to noise rather than signal. Where AI does deliver clear operational gains is inside structured enterprise architectures, where a 40% reduction in manual processing time has been documented. This review synthesizes these findings into a proposed algorithmic accountability architecture aligned with ISSB convergence efforts, and maps a realistic path toward democratising ESG compliance across firm sizes and geographies.

Index Terms—Algorithmic accountability, corporate financial performance (CFP), ESG rating standardization, ESG risk assessment, greenwashing detection, large language models (LLMs), machine learning, natural language processing (NLP), sustainable finance.

I. INTRODUCTION

It is worth pausing to consider just how much has changed in the relationship between investors and corporations over the past two decades. ESG criteria — once treated as optional add-ons to corporate social responsibility programmes — have become genuine drivers of capital allocation and corporate financial performance [1], [2]. The catalysts for this shift are well-documented: the 2006 UN Principles for Responsible Investment and the 2015 Paris Agreement both signalled, loudly and clearly, that non-financial risks could no longer be treated as footnotes to investment analysis [3], [4]. The market responded. ESG-aligned assets reached $35.3 trillion in 2020, with projections placing the figure above $50 trillion by 2025 [3], [5]. Firms are now assessed not just on quarterly earnings but on how they treat their workforce, manage natural resources, and structure their governance [2], [6]. The evidence that this matters financially is increasingly hard to dismiss — strong ESG performance has been linked to better Tobin's Q valuations, lower risk exposure, and a demonstrable buffer against reputational fallout when controversies do occur [4].

The problem is that the infrastructure supporting all of this has not kept pace. The ESG reporting landscape is genuinely fragmented in ways that create serious practical difficulties [1]. GRI, SASB, and TCFD all operate simultaneously, each with its own logic, without any binding hierarchy between them [1], [2]. Companies pick and choose what to disclose, and there is no authoritative standard against which to hold them. On top of this, the third-party rating agencies that investors rely on — Bloomberg, Thomson Reuters (LSEG), MSCI — arrive at meaningfully different scores for the same company using the same underlying information [3], [4]. That is not a minor calibration issue; it is a systemic problem that distorts how capital is allocated [2].

The voluntary, self-reported nature of most ESG data creates a further vulnerability that the literature describes as "greenwashing" — the practice of projecting a socially responsible image while the underlying operations tell a different story [2], [7]. Traditional financial reporting was not built to catch this, and the blind spots it creates are consequential [2]. There is also an equity dimension that

deserves more attention than it typically receives. The compliance costs associated with rigorous ESG reporting effectively exclude smaller businesses from sustainable finance markets [2], [8]. In developing economies, where micro-enterprises account for nearly 90% of industrial activity, most firms simply lack the resources to engage with complex ESG frameworks in any meaningful way [8].

AI, LLMs, and NLP have emerged as the most promising technical response to these challenges [9], [10]. Across the literature, researchers have deployed everything from simple regression models to BERT and LSTM networks to read corporate reports, detect misleading language, and generate ESG risk scores without relying on subjective third-party assessments [7], [10]. Enterprise platforms like SAP Analytics Cloud are already operationalising some of these capabilities, with documented efficiency gains [9]. But the technology also brings its own complications — algorithmic bias, data privacy concerns, and the persistent problem of model interpretability [9], [10].

This paper works through that landscape systematically. The goal is not simply to catalogue what has been tried, but to identify what actually works, where the serious gaps remain, and what a realistic accountability architecture might look like given the current state of both the technology and the regulatory environment [2], [7], [9], [10].

II. METHODOLOGY

This study follows a Systematic Literature Review (SLR) methodology structured around PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) guidelines. The intention was to be rigorous enough that the process could be replicated, and transparent enough that the selection decisions can be scrutinised.

A. Search Strategy

The search covered four databases: IEEE Xplore, Scopus, Web of Science, and Google Scholar. The Boolean search string was constructed to capture the intersection of ESG-domain terms ("ESG," "environmental social governance," "sustainable finance," "greenwashing"), computational methodology terms ("machine learning," "NLP," "artificial intelligence," "BERT," "LSTM," "deep learning"), and outcome terms ("risk assessment," "corporate financial performance," "ESG scoring"). The search was bounded to peer-reviewed journal articles, IEEE and ACM conference proceedings, and reputable institutional working papers published between 2020 and 2026.

B. Inclusion and Exclusion Criteria

Papers were included if they addressed ESG evaluation, rating, or reporting in some empirically grounded way — whether quantitative, computational, or rigorously qualitative. Papers were excluded if they were purely theoretical, if they addressed ESG only in the context of portfolio optimisation without any firm-level dimension, or if they were grey literature, editorials, or opinion pieces. The aim was to keep the sample anchored in what researchers have actually found, not what they expect or advocate.

C. Final Sample and Analysis Framework

After deduplication and quality screening, eleven papers were retained. They span panel OLS regression [4], [5], meta-analysis [3], NLP-based ML benchmarking [11], enterprise AI case studies [9], and qualitative thematic analysis [8], and they draw on data from international, Nigerian, and Indian firm-level contexts as well as S&P 500 companies. The analysis was organised around three thematic pillars: the ESG-CFP nexus; AI and NLP predictive modelling; and regulatory frameworks and greenwashing. Table I maps each paper against methodology, dataset scope, key findings, and limitations.

III. LITERATURE REVIEW

When you step back and look at what researchers have actually been doing in this space, the work falls fairly naturally into three areas: trying to measure whether ESG practices genuinely move the financial needle for companies; exploring whether AI and language models can do a better job of evaluating sustainability than human analysts; and grappling with the uncomfortable reality that the data feeding all of these efforts is still deeply unreliable.

A. The ESG-CFP Nexus: Econometric Approaches and Methodological Discrepancies

Much of the foundational work in this area takes a numbers-first approach, and some of it is genuinely compelling. Brighi et al. worked through data from over 10,000 international companies and found something that feels intuitively right: ESG controversies hurt firm value, but companies that had built up a strong ESG track record before a scandal broke fared considerably better than those that had not [4]. It is the corporate equivalent of having savings before a crisis. Yahaya reached a broadly similar conclusion looking at Nigerian listed firms, where strong ESG ratings appeared to signal credibility to investors in ways that translated into better market-to-book ratios and higher capital expenditure — a finding that matters especially because it comes from an emerging market context that often gets overlooked [5]. Koshy and Muthulakshmi's work on Indian companies added another layer, showing that even firms carrying heavy debt or posting operating losses could hold onto investor trust if their ESG reporting was consistent and transparent [6].

That said, Narula et al. raise a problem that quietly undermines much of this research [3]. When they ran a meta-analysis across existing empirical studies, they found something troubling: whether ESG appears to help or hurt financial performance depended heavily on which rating agency's data the researchers were using. Studies built on Bloomberg scores tended to find positive relationships, while those using Thomson Reuters data frequently found the opposite [3]. This is not a minor methodological footnote — it means that a significant portion of the econometric literature may be telling us more about the quirks of rating agencies than about ESG's actual effect on firm performance [3].

B. AI and NLP Interventions: The Performance Paradox in Predictive Modeling

Given the messiness of third-party rating data, it makes sense that researchers have started asking whether AI can cut through it — reading corporate reports directly and forming its own judgements. Ashok's study of AI-driven dashboards within SAP Analytics Cloud showed real practical promise: automated ESG reporting cut manual processing time by 40% and made audit preparation substantially easier [9]. Rane et al. similarly found that machine learning and remote sensing tools could meaningfully improve supply chain transparency and energy monitoring [10]. On the operational side of large enterprises, the technology works.

But here is where it gets interesting, and a little humbling for the deep learning community. When Nandiraju and Kanthi actually benchmarked several models — including LSTM networks and BERT — against simple Linear Regression for predicting ESG risk scores from real corporate annual reports, the sophisticated models lost [11]. Linear Regression hit 68.09% accuracy. LSTM managed only 25.53%, and BERT reached 46.81% [11]. The explanation the authors offer is convincing: corporate ESG reports are filled with repetitive, formulaic language that confuses models built to detect nuance and variation in text [11]. Simpler models, it turns out, are better at recognising broad patterns in uniform data. Beyond the model performance question, Kulkarni et al. flag a more human problem — smaller businesses simply cannot afford any of this infrastructure, and data privacy concerns and algorithmic bias risks compound the barrier further [8].

C. Regulatory Frameworks, Greenwashing, and the Standardization Deficit

All of the above — the econometric models, the AI pipelines, the enterprise dashboards — ultimately depend on the quality of what companies actually report. And that is where the literature identifies its most stubborn problem. Dako et al. make the point plainly: when GRI, SASB, and TCFD all coexist without any binding hierarchy, companies can and do choose to report what makes them look good [2]. That selective disclosure is precisely what greenwashing is, and it corrupts the data that every downstream model relies on. Biswas and Gupta add that even the ISSB's recent harmonisation push, while welcome, has not yet closed this gap [1].

Kalyanam's retail-sector research offers a partial way forward, showing that AI tools can cross-reference what companies claim against what their supply chain and operational data actually shows [7]. It is a promising direction. But as the literature makes repeatedly clear, until Scope 3 emissions data is consistently captured and regulators actually enforce disclosure standards rather than merely encouraging them [1], [2], [7], even the most sophisticated model is only as good as the corporate narrative it is fed — and right now, that narrative is still too easy to curate.

IV. COMPARATIVE ANALYSIS: ARCHITECTURES, ALGORITHMS, AND DATASET EFFICACY

Once you lay all eleven papers side by side and actually compare what each team did, a clear tension emerges — one that the field has not yet resolved. On one side, you have traditional econometric approaches that are methodologically rigorous but fundamentally dependent on whoever rated the ESG data in the first place. On the other, you have AI and NLP systems that promise objectivity but, at least in their current form, deliver it inconsistently. Table I maps this out in detail. What follows is an attempt to make sense of where these approaches agree, where they contradict each other, and what that means going forward.

TABLE I Systematic Cross-Paper Comparison of ESG Research Methodologies, Findings, and Limitations

Ref.	Authors (Year)	Methodology	Dataset / Scope	Key Findings	Limitations
[1]	Biswas & Gupta (2023)	Exploratory Review	Global ESG regulatory frameworks	Fragmented taxonomy impedes cross-portfolio benchmarking; ISSB unification efforts are nascent	No empirical validation; normative in scope
[2]	Dako et al. (2023)	Policy-Research Integration Model	ESG financial reporting frameworks	GRI/SASB/TCFD fragmentation enables selective disclosure and greenwashing	Framework-specific analysis; limited to reporting-side actors
[3]	Narula et al. (2024)	TCCM Framework + Meta-Analysis	Multi-ERP empirical studies	ERP choice (Bloomberg vs. Thomson Reuters) reverses ESG-CFP statistical direction	Dependent on quality of source studies; no primary data
[4]	Brighi et al. (2025)	Panel OLS; Firm/Year/Country FE; 78,738 obs.	10,212 international firms	ESG controversies reduce firm value; strong prior ESG buffers market fallout	Retrospective; ERP data not standardized across firms
[5]	Yahaya (2026)	Fixed-Effects Panel Regression; Hausman Test	148 Nigerian listed firms; 2,220 obs.	ESG ratings boost capital expenditure intensity and market-to-book ratios in emerging markets	Single-country scope limits generalizability
[6]	Koshy & Muthulakshmi (2023)	Financial Statement Analysis	Top Indian ESG-rated companies	ESG transparency sustains long-term investor loyalty despite short-term financial strain	Small sample; subjective ESG classification criteria
[7]	Kalyanam (2024)	Mixed-Methods; AI BI Tools	Retail sector ESG data	AI business intelligence effectively cross-references sustainability claims vs. circular economy outcomes	Sector-specific findings; limited scalability validation
[8]	Kulkarni et al. (2023)	Qualitative Thematic Analysis; NVivo	MSMEs across sectors	AI adoption severely constrained for SMEs by cost, data scarcity, and algorithmic bias	Qualitative only; no quantitative performance benchmarks
[9]	Ashok (2024)	Enterprise AI Case Study; SAP S/4HANA	SAP Analytics Cloud ESG module	40% reduction in manual ESG processing time; NLG improves GRI/ISSB compliance	Proprietary platform dependency; limited reproducibility
[10]	Rane et al. (2024)	Bibliometric Analysis; VOSviewer Clustering	ML & blockchain sustainability literature	AI and blockchain synergize for ESG optimization and supply chain traceability	Lacks empirical testing; conceptual in nature
[11]	Nandiraju & Kanthi (2024)	NLP/ML Benchmarking; ChatGPT-4 Pipeline	S&P 500 annual reports (web-scraped)	Linear Regression (68.09%) outperforms LSTM (25.53%) and BERT (46.81%) on ESG text	Limited to S&P 500 corpus; accuracy metric scope unclear

A. Econometric Frameworks and the Dataset Inconsistency Conflict

The econometric work in this space is, taken on its own terms, genuinely solid. Brighi et al. worked with nearly 79,000 firm-year observations and used firm, year, and country fixed effects precisely to rule out the kind of background noise that can make results look more meaningful than they are [4]. Yahaya's study of Nigerian listed companies was similarly careful — the Hausman test result ($\chi^2 = 87.42$) pointed clearly toward the fixed effects model, which matters because it tells us that what drives ESG outcomes is tied to something specific about individual firms, not just general market conditions [5]. Both studies point in the same direction: companies that take ESG seriously tend to score better on Tobin's Q and invest more heavily in capital expenditure [4], [5].

The problem surfaces when Narula et al. zoom out and look across the literature rather than at any single study [3]. What they found is genuinely uncomfortable: run the same analysis with Bloomberg ESG data and you tend to get a positive relationship between ESG and financial performance. Switch to Thomson Reuters data for the same companies and the same metrics, and the relationship flips or disappears entirely [3]. That is not a small discrepancy you can footnote away. It means that a significant portion of the econometric literature may essentially be reporting the preferences of rating agencies rather than any underlying truth about ESG's financial impact. Until there is a common standard for how ESG is measured in the first place, building elaborate regression models on top of that data is like constructing a precise measurement on a foundation that shifts depending on who poured the concrete [1], [3].

B. Algorithmic Architectures: The Deep Learning Performance Paradox

This is probably the most counterintuitive finding in the entire body of literature, and it deserves more attention than it typically receives. Nandiraju and Kanthi built a genuinely sophisticated pipeline — using ChatGPT-4 to parse annual reports directly from S&P 500 companies, then benchmarking four different model types against each other [11]. On paper, LSTM networks and BERT should have dominated. LSTM is designed to capture long-range dependencies in sequential text, and BERT reads language in both directions simultaneously, giving it richer contextual understanding than almost any other architecture available. And yet, when the results came in, Linear Regression outperformed both of them — 68.09% accuracy against LSTM's 25.53% and BERT's 46.81% [11].

The explanation the authors offer is worth sitting with. Corporate ESG reports are not written the way most text is written. They are formulaic, repetitive, and carefully managed — every company essentially saying the same things in slightly different arrangements of the same words [11]. Deep learning models are built to detect subtle variation and nuance in language; that is precisely what makes them powerful in most NLP contexts. But when the text is structurally homogeneous by design, those same capabilities become a liability. The model chases linguistic patterns that do not carry real signal, while a simpler linear model just picks up the broad statistical trends and does a better job of it [11]. Model complexity and model usefulness are not the same thing, and domain characteristics matter enormously when choosing an approach.

C. Enterprise Scalability and Algorithmic Integration

Where AI does clearly work well is inside structured enterprise environments, and Ashok's case study makes this point persuasively [9]. The SAP Analytics Cloud implementation he evaluated was not trying to read raw corporate prose and infer ESG scores from it. Instead, it pulled structured data from IoT emissions sensors, ran anomaly detection algorithms on that data, and used Natural Language Generation to produce compliance-ready reports in GRI and ISSB formats automatically [9]. The result was a 40% reduction in manual processing time and meaningfully better audit preparedness [9]. The lesson is not that AI cannot handle ESG — it is that AI handles ESG much better when it has clean, structured data to work with rather than the polished but unreliable narratives companies publish.

Rane et al.'s bibliometric mapping work using VOSviewer points toward where the field is heading conceptually [10]. But Kulkarni et al. bring the conversation back to earth with an important reality check [8]. The firms that could most benefit from better ESG tooling — MSMEs in developing markets — are almost entirely locked out of these systems. The costs are prohibitive, the technical requirements are demanding, and the algorithmic bias risks are real and underexplored [8]. Building enterprise AI for the S&P 500 is one thing; building it for a mid-sized manufacturer in a developing economy is an entirely different problem that the literature has barely begun to address.

D. Synthesis of Findings

If there is one thing all eleven papers agree on, it is that paying attention to ESG improves outcomes — financially, operationally, and in terms of risk management [1]–[11]. That much is settled. Everything else is considerably more contested. The deep learning models that dominate headlines are currently struggling in this domain, and simpler approaches are outperforming them on the tasks that matter most [4], [5], [11]. Enterprise AI works, but only for organisations large enough to afford it and only when the underlying data is structured [9]. And critically, none of it functions reliably until the data it feeds on is standardised. The most important work right now may not be algorithmic at all, but regulatory. Until TCFD, SASB, and GRI align around a common disclosure taxonomy, hybrid systems — pairing straightforward predictive models with real-time operational data — represent the most defensible approach currently available [1], [9].

V. RESULTS AND DISCUSSION

Stepping back from the individual papers and asking what they collectively tell us produces four findings that are significant enough to shape how the field thinks and invests going forward. Each one has direct implications for practitioners trying to build better ESG systems, for

policymakers trying to regulate them, and for researchers deciding where to focus next.

A. The ERP Standardization Imperative

The most structurally important finding of this review is also the most uncomfortable one: no amount of algorithmic sophistication can compensate for unreliable input data. Narula et al.'s evidence that simply switching from Bloomberg to Thomson Reuters data reverses the measured relationship between ESG and financial performance [3] means that any ML model trained on non-harmonised ERP data will inherit those biases silently. The model will be internally consistent and externally misleading. This finding makes the ISSB's mandate for GRI, SASB, and TCFD convergence [1], [2] not just a regulatory nicety but a technical prerequisite. The field is currently building increasingly sophisticated models on a foundation that is not stable, and that is a problem no amount of better architecture will solve.

B. Rethinking the Deep Learning Paradigm for ESG NLP

The performance paradox identified by Nandiraju and Kanthi [11] challenges something the NLP community has largely taken for granted — that more powerful models produce better results. In the ESG domain, that assumption breaks down because corporate sustainability disclosures are structurally homogeneous in a way that most NLP training data is not. BERT's bidirectional attention and LSTM's sequential memory are liabilities when every document uses the same phrases arranged in the same order. What this points toward, practically, is the need for domain-adapted pre-training on ESG-specific corpora — something analogous to what FinBERT did for financial text. Until that exists, researchers and practitioners should resist the assumption that state-of-the-art NLP automatically translates into state-of-the-art ESG risk scoring.

C. Hybrid Enterprise Architectures as the Near-Term Optimum

The 40% efficiency gain Ashok documented [9] is the most operationally actionable finding in this review, and it is worth understanding why the architecture produced it. The system was not trying to interpret corporate prose — it was processing structured IoT sensor data, running anomaly detection against that ground truth, and generating compliance reports from the results. By bypassing self-reported text entirely in favour of verifiable operational data, it sidestepped the three main vulnerabilities the rest of the literature keeps running into: unreliable input data, greenwashing susceptibility, and the complexity costs of advanced NLP. Hybrid architectures of this kind currently represent the nearest thing to a working solution that the evidence supports.

D. The MSME Democratization Gap

Kulkarni et al.'s qualitative findings [8] introduce an equity dimension that the technically-focused literature largely ignores. The firms that would benefit most from better ESG assessment tools are not the S&P 500 companies that researchers tend to study — they are the MSMEs that make up the vast majority of global economic activity but are effectively priced out of sustainable finance markets because compliance costs are prohibitive [8]. There is a self-reinforcing dynamic at work here: MSMEs cannot afford ESG compliance, so they cannot access ESG-linked capital, so they cannot modernise. Addressing this requires purpose-built, cloud-native tools designed around MSME resource constraints rather than enterprise budgets, backed by regulatory incentives that acknowledge the proportionality problem the ISSB framework gestures toward but has not yet fully resolved [1], [2].

VI. CONCLUSION

This review set out to make sense of a literature that is growing rapidly, pulling in multiple directions, and producing results that sometimes contradict each other. Across eleven studies spanning econometric modelling, NLP benchmarking, enterprise AI deployment, and qualitative analysis, four conclusions emerge clearly enough to be stated with some confidence.

First, the reliability of any ESG risk model depends more on the quality and consistency of its input data than on the sophistication of its architecture. Narula et al.'s finding that ERP choice alone reverses the measured ESG-CFP relationship [3] makes global regulatory harmonisation — specifically, genuine convergence of GRI, SASB, TCFD, and ISSB taxonomies [1], [2] — a necessary precondition for scalable algorithmic accountability, not just a governance aspiration.

Second, the field needs to take the performance paradox seriously. Linear Regression outperforming LSTM and BERT by a substantial margin on ESG text [11] is not a fluke — it reflects something real about the nature of corporate sustainability disclosures that current deep learning architectures are poorly equipped to handle. Domain-adapted transformer models, trained on ESG-specific corpora, are the most logical next step.

Third, hybrid enterprise architectures that combine structured operational data with AI-driven anomaly detection currently represent the optimal near-term solution [9]. They deliver documented efficiency gains precisely because they bypass the weaknesses of self-reported text rather than trying to compensate for them algorithmically.

Fourth, the MSME democratisation gap is a structural problem, not a marginal one. Until the tools that work at enterprise scale become accessible at MSME scale, sustainable finance will remain concentrated among the firms that need it least [8].

The principal limitation of this review is the sample size of eleven papers, which, while methodologically diverse, cannot claim comprehensive coverage. Several studies also operate in single-country contexts — Nigeria [5] and India [6] — whose findings may not transfer directly to other markets. Future work should prioritise five directions: developing ESG-domain-adapted transformer models; running longitudinal studies that track ML performance as ISSB-harmonised data becomes available; designing and piloting low-cost MSME ESG assessment platforms; investigating federated learning as a privacy-preserving architecture for cross-border ESG data

sharing; and building explainable AI (XAI) frameworks that give regulators and investors genuine insight into how algorithmic ESG scores are produced [9], [10].

REFERENCES

[1] B. Biswas and R. P. Gupta, "Global ESG Regulatory Framework and Sustainability: Issues and Prospects," PG Department of Commerce, University of Calcutta, 2023.

[2]O. F. Dako, T. A. Onalaja, P. S. Nwachukwu, and F. A. Bankole, "Integrating ESG Performance Metrics into Financial Reporting Frameworks to Strengthen Sustainable Investment Decision-Making Processes," Int. J. Adv. Multidisc. Res. Stud., vol. 3, no. 2, pp. 1239–1252, 2023.

[3]R. Narula, P. Rao, S. Kumar, and A. Paltrinieri, "ESG investing & firm performance: Retrospections of past & reflections of future," Corp. Social Responsib. Environ. Manag., 2024.

[4]P. Brighi, A. C. F. Della Bina, and V. Venturelli, "Firm value and risk: how relevant are ESG factors and ESG controversies?" J. Financ. Rep. Account., 2025.

[5] O. A. Yahaya, "ESG ratings and investment decisions in Nigeria,"
J. Manage. Finance, vol. 15, no. 2, pp. 51–78, 2026.

[6] B. A. Koshy and M. R. Muthulakshmi, "An empirical study on the financial performance of top sustainable companies in India in the ESG sector," Int. J. All Res. Educ. Sci. Methods, vol. 11, no. 12, Dec. 2023.

[7] A. K. Kalyanam, "Integration of ESG data in the retail sector: Essential for sustainable business and risk management," Int. J. Sci. Res. Eng. Manag., vol. 8, no. 12, Dec. 2024.

[8]A. Kulkarni, S. Joseph, and K. Patil, "Role of Artificial Intelligence in sustainability reporting by leveraging ESG theory into action," in Proc. 2023 Int. Conf. Advancement in Computation & Computer Technologies (InCACCT), Pune, India, 2023.

[9] P. P. K. Ashok, "AI-driven ESG dashboards within SAP Analytics Cloud for enterprise reporting," Eur. J. Adv. Eng. Technol., vol. 11, no. 5, pp. 157–162, 2024.

[10] N. L. Rane, S. P. Choudhary, and J. Rane, "Artificial intelligence driven approaches to strengthening Environmental, Social, and Governance (ESG) criteria in sustainable business practices: A review," University of Mumbai, India, 2024.

[11]A. Nandiraju and S. Kanthi, "Leveraging multimodal large language models and natural language processing techniques for comprehensive ESG risk score prediction," in Proc. 6th Int. Conf. Finance, Economics, Management and IT Business (FEMIB), 2024.